



Issue 19, April 28, 1999

Feature Story

Each month we provide a feature article on key industry trends and developments. Authored by a member of Intel's Executive Staff, it offers insightful and useful information for product development, planning and execution.

Inside Looking In

Senior Technical Marketing Manager for Platform Technologies Tim Mostad gives you a fresh perspective on the latest technologies making their way onto Intel Architecture platforms. Tim lets you see the work through the eyes of the people making it happen and lets you hear what they really think. It's straight talk from developers to developers.

Pentium® III Processor Platform Series

Pentium® III Processor Platform Series PSN will bring you focused articles on the hottest technologies about to arrive on Pentium III processor-based platforms. From Rambus memory technology to optimizing software, it's news and information you can definitely use, and you'll find it every month here in PSN.

Top News Stories

Delivering an in-depth report on key platforms, products and technologies, our Top Stories provide a monthly source of information on the issues affecting hardware developers. Be sure to check in every month for the latest stories that are driving the evolution of the industry.

Reader Services

If you are new to [Platform Solutions](#) and would like to receive this companion newsletter to the on-line version, please visit [Platform Solutions](#) on-line and go the "[Subscribe Now](#)" section to register and sign up for delivery. The on-line version provides lots of direct links for quick access to the developer information and news reported in each issue, whether it's on Intel's Web site or industry Web sites.

We want you to consider *Platform Solutions* as your personal information resource for the Intel® Architecture platform. If you can help us make it better, or if you have a comment, question or a specific topic you would like to see covered, we want to hear from you. Please take the opportunity to send us an e-mail with your specific feedback or request to: platform.solutions@intel.com

If you do not want to receive this mailing in the future, please send an email to: platform.solutions@intel.com with "unsubscribe" in the body of the message.

On behalf of all of us at Platform Solutions, welcome to the future of the PC platform!

DISCLAIMER: THE MATERIALS ARE PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND INCLUDING WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT OF INTELLECTUAL PROPERTY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT SHALL INTEL OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION, LOSS OF INFORMATION) ARISING OUT OF THE USE OF OR INABILITY TO USE THE MATERIALS, EVEN IF INTEL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME JURISDICTIONS PROHIBIT THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT APPLY TO YOU. INTEL FURTHER DOES NOT WARRANT THE ACCURACY OR COMPLETENESS OF THE INFORMATION, TEXT, GRAPHICS, LINKS OR OTHER ITEMS CONTAINED WITHIN THESE MATERIALS. INTEL MAY MAKE CHANGES TO THESE MATERIALS, OR TO THE PRODUCTS DESCRIBED THEREIN, AT ANY TIME WITHOUT NOTICE. INTEL MAKES NO COMMITMENT TO UPDATE THE MATERIALS.

Table of Contents

[\(Click on descriptions to jump to articles\)](#)

FEATURE.....	3
INTEL BRINGS NEXT-GENERATION TECHNOLOGIES TO VALUE PCs.....	3
INSIDE LOOKING IN.....	5
WHEN “TOO GOOD” IS NOT GOOD ENOUGH.....	5
FOCUS.....	7
INTEL® GRAPHICS TECHNOLOGY: YOUR WINDOW TO THE WEB	7
PENTIUM® III PROCESSOR PLATFORM SERIES.....	11
USING STREAMING SIMD EXTENSIONS TO BOOST SPEECH RECOGNITION PERFORMANCE	11
TOP STORIES	14
INTEL® 752 GRAPHICS ACCELERATOR: THE COMPLETE MULTIMEDIA SOLUTION	14
NETWORK PROCESSORS HELP ENABLE THE NEW INTERNET ECONOMY.....	16
SELF-TEST TOOL ACCELERATES WFM 2.0 IMPLEMENTATION.....	19
THE BENEFITS OF INTEGRATED RAID	21
POWER BUDGETING SOFTWARE PROTECTS THE INSTANTLY AVAILABLE PC	23
TECHNOLOGY NEWS BYTES	25
INDUSTRY EVENTS	28
MAY 1999.....	28
JUNE 1999.....	29

Feature

Intel Brings Next-Generation Technologies to Value PCs

By Patrick P. Gelsinger
Vice President and General Manager
Desktop Products Group
Intel Corporation

PC consumers are looking for great performance and exceptional value. Intel is helping the industry meet the challenge with next-generation technologies for Value PCs.

Pick up just about any popular consumer computer magazine these days, and chances are the cover will feature a dramatic headline on one or more of these topics:

- How to get a PC with great performance.
- How to get on the Internet.
- And, how to get all of this for less than \$1,000.

While new technologies for mainstream and enthusiast PCs continue to grab their share of the headlines, the low-cost Value PC market segment continues to attract increasing attention from consumers, many of whom are buying a PC for the first time.

These market segment dynamics pose an interesting challenge for the industry. How can we build PCs that provide great performance while offering exceptional value? While it is true that yesterday's platform technologies can sometimes be adapted to achieve lower costs, this development model often fails to provide the quality experiences that customers expect from today's Internet-based and graphics-intensive applications.

Giving consumers more for their money requires next-generation technologies, and Intel is helping the industry meet that challenge.

Intel® Celeron™ Processor 466 MHz

We are introducing the Intel® Celeron™ processor at 466 MHz, the fastest processor Intel has ever developed for Value PCs. It delivers the performance needed to support customer satisfaction with popular Internet applications, educational titles, interactive 3D games, and productivity programs. While it brings high performance, the Celeron processor also provides the industry with leading-edge Intel technology, supported by the Intel brand quality and reliability that is recognized worldwide.

Cost reduction is a particularly important issue in this segment, and the new PPGA packaging of the Celeron processor enables developers of Value PCs to markedly reduce overall platform costs.

Intel® 810 Chipset

The Intel® 810 chipset is Intel's first-ever 3D graphics integrated chipset featuring Intel® Scalable Graphics Architecture. It delivers solid Accelerated Graphics Port (AGP) performance, while enabling cost reductions through the smart integration of next-generation 3D graphics and bus technologies:

- **Intel Direct AGP and Dynamic Video Memory** delivers AGP-class graphics and compatibility, while saving silicon, pin and bus-routing costs. Intelligent arbitration of video and memory resources optimizes throughput and graphics performance.
- **Intel Accelerated Hub Architecture** eliminates the PCI bus bottleneck, significantly improving internal bandwidth for improved concurrency and more lifelike audio and video.

Intel® Drivers

Intel has made a major commitment to the development of world-class quality drivers for processors, chipsets, and graphics accelerators. More than 200 Intel software engineers are working to support developers with drivers that offer performance, robustness and compatibility, while helping developers meet their time-to-market goals. The application of best-known chipset validation methodology to drivers allows Intel's driver division to complement the launch of new products with robust drivers that support a range of new technologies, including "soft DVD," AC'97 soft audio and soft modem technology, flat-panel display technology and video capture. Intel is also working to develop a common driver model across multiple product lines that will allow developers to consolidate their validation effort and bring multiple platforms to market simultaneously.

Integrated Technologies

The Intel 810 chipset also integrates a number of other leading-edge technologies for balanced platform performance and reduced development costs. Integrated features include:

- A Digital Video-Out (DVO) port which enables users to connect to the latest digital TV-Out and digital flat-panel display components
- AC'97 controller for soft audio/modem technology capability
- Hardware motion compensation for host-based MPEG-2* decode
- Support for ultra-fast ATA-66 hard disk drives
- Instantly Available PC technology
- Direct support for multi-monitor technology

Moving in a New Direction

Intel technology is helping to move the Value PC market segment in an exciting new direction in 1999. With smart integration that brings next-generation technologies to Value PCs, the Intel Celeron processor and Intel 810 chipset are delivering on the promise of great PC performance and a quality Internet experience—at an exceptional value.

For More Information

Visit the Intel Developer Web site for information on the [Intel Celeron processor—466 MHz product overview](#) and the [Intel 810 Chipset product overview and technical documentation](#)

Read this month's focus story on [Intel® Graphics Technology](#): Your Window to the Web.

About the Author

Mr. Pat Gelsinger is the vice president and general manager of Intel Corporation's Desktop Products Group. Mr. Gelsinger joined Intel in 1979 and has 20 years of experience in general management and product development positions.

He holds six patents, more than a dozen applications, as well as more than 20 publications in the high-tech field. His promotion to group vice president at age 32 made him the youngest vice president in the history of Intel.

Inside Looking In

When “Too Good” is Not Good Enough

By Tim Mostad
Senior Technical Marketing Manager
Intel Corporation

While plugfests serve an important function in the early to middle stages of technology development, we need a new testing paradigm for the long term.

The goal of my technical marketing organization is to enable the industry to adopt and implement a continuous stream of new platform technologies rapidly and efficiently. One of the tools we use is nicknamed a Plugfest (also known as Interops, Compliance Workshops, BootCamps, and various other names depending upon the audience). The simplicity of the concept is its greatest value. Products need to inter-operate before they ship to end users. In the later stages of the development process, we invite system and add-in makers to try their products together. Since 1992, when we invented them to expedite the launch of PCI, Plugfests have been effective tools for getting the next new capability bootstrapped.

At Intel, I've managed to get entirely too much credit for inventing Plugfests. Actually, the genesis of the idea was a conversation between me and Mike Bailey, my boss at the time. We were concerned about how we'd get members of the PCI Special Interest Group to rent space in our Comdex suite. We needed to provide public proof of PCI critical mass so it would result in immediate acceptance and instantaneous explosive growth. (Ok, so that's a *slight* exaggeration. But it's nearly what happened and since I'm getting too much credit anyhow, a little more can't hurt.) We also realized companies were afraid to show untried products. So, on short notice, we invited PCI SIG companies together for a "PCI Compliance Workshop," the first-ever Plugfest. There was good turnout and consequently the Comdex suite was a hit too. As they say, "the rest is history." In reality, Mike and I just applied an old idea in a novel way. Since we know that what gets measured gets fixed, measure interoperability at a Plugfest and it will get fixed.

The concept was simple enough but we immediately knew we had an implementation hurdle. How do you both test one brand new, potentially broken product with yet another brand new, potentially broken product *and* continue to improve testing as the products improved so that everything continued to get better? It would be nice if the development process was like a phonograph record, get the industry needle in the groove and keep it spinning. This would get you smoothly from technology start to end. Instead, technology development is more like a CD ROM player. You complete a cycle, and then you have to perform a step function to make the next cycle worthwhile.

We had to get the process started so the tests we used at the first event were rudimentary, and passing them wasn't very difficult. We took what we learned and improved the testing at the next event. Products and tests would improve in lock step until we had tests that would come as close as possible to guaranteeing conformance to the PCI Interface Specification and we had products that actually worked together. Assuming a well-written specification, the problem was solved. Or was it?

My "PCI life" theoretically ended long ago but it is tough to be far from PCI if you are in the PC technology business. It's simply become ubiquitous. It's also recently become the most maddeningly frustrating time-wasting pain I could imagine.

Probably as a function of my job, my home PC is a never-ending project. I'm always trying to stuff more stuff in the proverbial 5 pound bag. Most recently, the bag was an Intel® Pentium® III processor, Intel® 440BX AGPset, one AGP/five slot system board with all slots except the one ISA slot full. Naturally, some of the stuff I was stuffing was PCI sound and modem cards. I figured the low latency of devices on the PCI bus would eliminate long waits for ISA transactions and unleash the performance of the processor. And the now-mature PCI technology would make this a breeze. Terrific idea, but something went wrong along the way.

According to the well-written PCI Interface Specification, all devices on the PCI bus must be able to share interrupts. This approach compensates for the legacy interrupt limitation of the PC/AT architecture. Highly optimized interrupt handling both in hardware and software make interrupt sharing possible with minimal performance impact. At the same time, it eliminates the problem of struggling with hardware interrupt jumpers. Moreover, USB successfully employs interrupt sharing by design. This has got to work, right? NOT. None of my add-in PCI devices would tolerate sharing an interrupt with another.

These were not your no-name, bargain basement products, either. They were the household names in the add-in market, well known even outside my household. The two different modems were the fussiest, failing to do anything when sharing an interrupt with another device. The sound card worked for just a second, until the other device generated an interrupt, then it seized up tighter than a door hinge on the Titanic. Even my trustworthy, fast and reliable SCSI card generated random disk errors when sharing its interrupt. And furthermore, they had all been through the Plugfest process. What was going on here?

I'm afraid we've architected a process that is both too good and not good enough. The continuing PCI Compliance Workshops—something like 19 have occurred—have given their attendees a lot of confidence in their products and the technology, apparently too much. We only ever intended Plugfests to serve as a technology launch vehicle, a purpose that they serve and will continue to serve perfectly. However, after we believe a technology is successfully launched in the industry, we naturally disengage and move on. We expected that Plugfests would experience a similar life cycle but they seem to cycle on and on endlessly, about as useful as a CD stuck playing the same track over and over. It's clear that products and technologies need continual testing, but an everlasting Plugfest is not the answer.

In short, we need a new long-term testing paradigm. We need one that matches open industry processes with the motivation to focus on testing rigor. A glimmer of what's possible exists in the System Test Implementers Forum (ST-IF). Right now, the organization's main objective is to pull together a comprehensive test suite in support of the PCXX efforts co-sponsored by Intel and Microsoft. However, it doesn't take a huge leap to imagine the establishment of a possibly third-party testing process in conjunction with the ST-IF—the rigorous kind of testing we'll need if technologies like PCI are to achieve our greatest expectations.

At Intel, we're beginning to kick this idea around and we're very interested in industry input and participation in the outcome. Until we get something going, tune in to the [System Test IF process](#) and get involved.

In the meantime, even without the long-term compliance testing problem fixed, my PC does run, but it requires setting each PCI device on its own interrupt. It also requires disabling a thing or two, putting certain devices in the "magic" slots, and setting the BIOS just right. (If you're interested in knowing more, send me a message at tim.mostad@intel.com. Now that my PC is working, it will continue to work at least until that next new technology I can't live without appears. Then, once again, I'll be running my own personal Plugfest. For the sake of having a life after PCI we need to get this compliance testing process fixed—and soon.

About the Author

Tim Mostad says, "the majority of my 18 years at Intel have been spent in the pursuit of technical marketing nirvana." He is responsible for demos, white papers, plugfests, and technical training to support the adoption of new desktop technologies.

Focus

Intel® Graphics Technology: Your Window to the Web

By Gary Thomas
General Manager
Graphics Components Division
Intel Corporation

Intel® Scalable Graphics Architecture will ultimately integrate Internet, video, 3D, and display technologies. Intel's graphics technology vision also includes graphics engines capable of higher levels of texture compositing and more operations per pixel.

Advances in Visual Computing are driving the PC graphics pipeline to the threshold of a new era. For example, Internet applications will soon add video compositing and 3D video effects. Video on the PC will evolve from today's MPEG-2* decode to PC/VCR technology and, ultimately, to HDTV. The domain of 3D graphics will add GUI compositing and curved surfaces. Higher display resolutions are emerging. All of these trends are placing increasing demands on the graphics solution, making it necessary to minimize bandwidth bottlenecks and increase the effective number of operations per pixel.

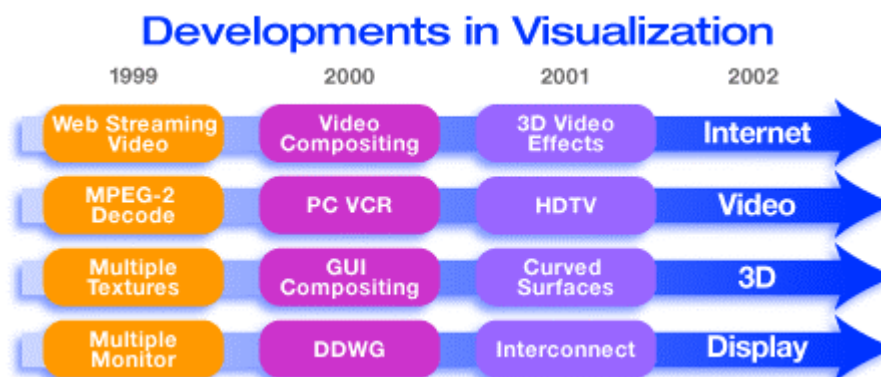


Figure 1. Developments in Visualization

The problem is that as the graphics pipeline evolves to meet these requirements, compatibility can become a matter of real concern for MIS departments that need to maintain longevity of software over PC hardware transitions. The current software driver model for graphics does not meet this requirement, which makes a separate validation and integration cycle necessary for each new platform.

Intel® Graphics Technology

Intel's long-term vision involves the development of superior graphics engines, beginning with the introduction of the Intel® 752 graphics accelerator and the Intel® 810 chipset, that can address the need for higher levels of texture compositing and more operations per pixel. The Intel® Graphics Technology roadmap will include ongoing advances in bandwidth to meet these requirements. In addition, as higher display resolutions emerge, typical of the digital broadcast standards, bandwidth between the PC and display will also increase. These requirements must also be balanced on bandwidth demands. These technology trends will move graphics solutions to multiple levels of compositing, while increasing the number of operations that can be achieved per pixel.

New directions in the software driver model have also emerged. Graphics device class driver models built within the Windows* operating system will support higher compatibility over PC transitions. On the hardware side, a move toward a true graphics instruction set architecture will significantly improve the stability of Visual Computing.

Intel® Architecture processors are designed to meet the requirements of specific market segments, and Intel Graphics Technology will complement these segments with graphics solutions that represent focused implementations to enhance Visual Computing.

Emergence of Intel® Scalable Graphics Architecture

Intel® Scalable Graphics Architecture will ultimately integrate Internet, video, 3D, and display technologies. Advances in performance will be needed as 3D texturing becomes the foundation of desktop PC graphics (2D operation will become a subset). Performance will also need to keep step as texture and fill rates continue to increase and highly realistic pixel blend techniques emerge in mainstream PC applications. The result of Intel's focused implementation strategy across all PC segments will be cost-sensitive, high-performance solutions and higher levels of integration, principally driven by the economic requirements of the Value PC segment.

Graphics Core Architecture

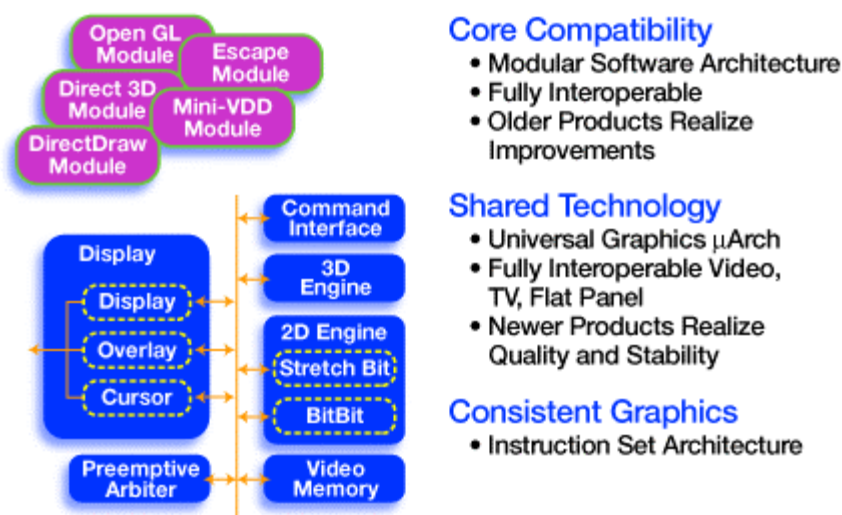


Figure 2. Graphics Core Architecture

Graphics Compatibility

Intel has invested significantly to develop the core competencies, graphics technology, and leadership in visual quality needed to support this effort. The highly successful introduction of the Intel740™ graphics accelerator was an example of this effort. With the introduction of the [Intel 752 graphics accelerator](#) and the [Intel 810 chipset](#), Intel is launching the first phase of a broad segmentation strategy that emphasizes platform solutions and software stability.

The next step will be to implement a Graphics Core Architecture with these characteristics:

- Core compatibility achieved through a modular software architecture that is fully interoperable, ensuring that older products can share the latest improvements.
- Shared technology based on a universal graphics micro-architecture that supports fully interoperable video, TV, and flat-panel display functionality, enabling newer products to achieve quality and stability.
- Consistent graphics based on instruction set architecture.

Scalable Graphics Performance

Scalability over time is an issue of great importance to MIS managers who must focus on eliminating potential system instability and reducing the cost of maintaining and upgrading PCs. Scalable graphics performance enables MIS departments to effectively plan and migrate software, while preserving compatibility.

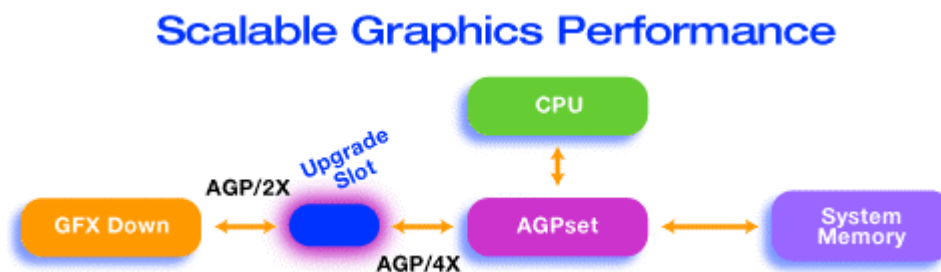


Figure 3. Scalable Graphics Performance

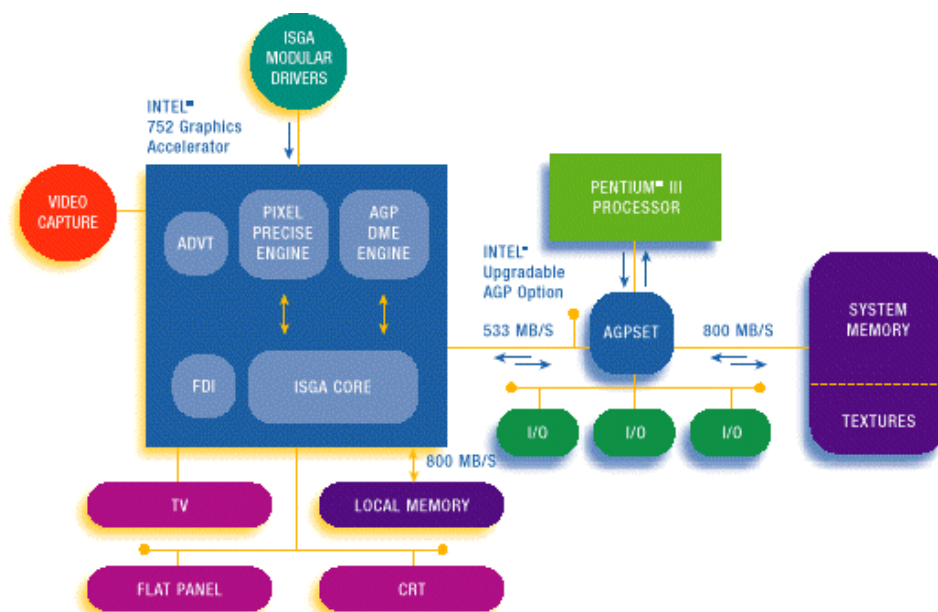
Intel Graphics Technology will achieve scalable graphics performance in three ways:

- Synchronizing transactions between the CPU and the graphics component.
- Providing scalable cores that allow modifications to frequency, pipeline width, and memory technologies.
- Implementing Scalable AGP technology, including Intel® Direct AGP and Intel® Upgradeable AGP.

Intel's goal is to support forward and backward software compatibility that will provide faster time-to-market for developers by enabling debugging with production-class drivers. Software compatibility will also enable safer technology migration for MIS departments, by supporting a familiar software image on newer hardware. OEMs will be able to reduce development costs by amortizing software customization over multiple platforms.

Moving to Next-generation Visual Processing and AGP4X

Continuing enhancements to Intel's Pixel Precise Engine will optimize visual quality and performance. On-going improvements will be made to the scalable graphics core with graphics instruction, the sub-pixel precise accuracy of the graphics engine, in addition to the hyperpipelined architecture of the graphics controller.



Intel® 752 Graphics Accelerator Block Diagram

For example, the newly introduced Intel 752 graphics accelerator features a Pixel Precise Engine with 16 tap anisotropic filter, emboss bump mapping, texture compression, and texture compositing. The forthcoming Intel® 754 graphics accelerator with Intel Upgradeable AGP technology is driver- and pin-compatible with the Intel 752 component and allows developers to easily enhance performance and cost effectively add AGP4X technology.

The Graphics Foundation for the Next Generation

Intel Graphics Technology is launching graphics products for all PC market segments, implementing a graphics architecture that delivers scalable performance, stability, and compatibility for Visual Computing. The industry should look to Intel for continuing breakthroughs in visual quality, delivered in scalable solutions with Intel quality, capacity, validation, software support, and worldwide brand recognition.

For More Information

Visit the Intel Developer Web site for more details on the [Intel 752 Graphics Accelerator](#) and the [Intel 810 Chipset](#).

About the Author

Gary Thomas is the general manager of the Graphics Components Division in Intel's Computing Enhancement Group, responsible for the development and marketing of graphics products.

Since joining Intel in 1979 as an engineer in the Telecommunication Products Operation, Thomas has held a variety of management positions at Intel. In 1993, he was promoted to the director of engineering for the PCI Components Division, responsible for the architecture, design, and manufacturing ramp of Intel® PCIsets resulting in the establishment and growth of Intel's chipset business. In 1995 he moved to Santa Clara to head Intel's 64-bit architecture microprocessor development. From 1996 to 1998, Gary was the general manager of the Santa Clara Processor Division, responsible for the development of IA-64 microprocessor products.

Gary Thomas received a B.S. in electrical engineering from the University of Illinois in 1977, and an M.S. in electrical engineering from Stanford University in 1979.

Pentium® III Processor Platform Series

This new column provides the latest information on platform technologies arriving with the Intel® Pentium® III processor.

Using Streaming SIMD Extensions to Boost Speech Recognition Performance

By Antonio C. Valles
Software Development Engineer
Software Performance Lab
Intel Corporation

Intel® Pentium® III processor Streaming SIMD extensions can provide a significant performance gain in speech recognition applications, for faster recognition and increased accuracy.

Speech recognition provides a very exciting and useful interface in communicating with the world of computers and the Internet [Quinnel95]. This technology has attracted the interest of both businesses and consumers. Yet, as applications for this technology grow, the demands for improved speech recognition capabilities have increased.

The most common statistical model used in speech recognition systems is the Hidden Markov Model (HMM). An HMM consists of a Markov chain of n states and three probability functions (see fig. 1). HMMs can be used to model a speech sample through a set of state movements governed by the three probability functions.

In a given HMM, many different state sequences can describe the speech sample, but each state sequence is associated with a different probability. The state sequence with the best probability is the path that can best describe the speech sample. This state sequence is known as the "best path." Doing an exhaustive search on an HMM for the best path is impractical. Instead, the Viterbi decoding algorithm offers a more practical approach. The Viterbi decoding algorithm is a time-intensive algorithm that can be optimized on an Intel® Pentium® III processor to increase speech recognition performance.




Figure 1. The parameters of a discrete, jump constrained Hidden Markov Model

With the introduction of the Streaming Single Instruction Multiple Data (SIMD) Extensions in the Pentium III processor, complex algorithms such as the Viterbi decoding algorithm can now be accelerated through the use of SIMD programming techniques. While this requires that programmers learn new instructions and identify how to best apply them to specific software applications, the performance improvements can be significant enough to compensate for the initial learning curve.

How Streaming SIMD Extensions Work

Traditional integer operations utilize a single stream of data; Pentium III SIMD instructions take advantage of the symmetry in algorithms to process multiple data simultaneously. The Streaming SIMD Extensions consist of approximately 70 new SIMD floating point, SIMD integer, and memory management instructions. These instructions can bring significant performance improvements to speech recognition applications.

Applying Streaming SIMD Extensions to Speech Recognition

There are three steps to the Viterbi decoding algorithm [Rabner89]. The step that contains most of the processing time is the recursion step. The recursion step can be expanded as follows:

$$\text{Dist}(j,t+1) = \min \left[\{ \text{Dist}(j,t) + a\text{Prob}(j,j) \}, \{ \text{Dist}(j-1,t) + a\text{Prob}(j,j-1) \}, \{ \text{Dist}(j-2,t) + a\text{Prob}(j,j-2) \} \right] + b\text{Prob}(j,k(t+1)) \quad \begin{matrix} 1 \leq t \leq T \\ 1 \leq j \leq N \end{matrix}$$

Figure 2. The recursion step of the Viterbi decoding algorithm

Here is how the equation looks when placed in SIMD form. Notice that the left column is the scalar version of the equation given above.

SIMD (4 Distances computed at once)			
Dist(j,t+1)	Dist(j-1,t+1)	Dist(j-2,t+1)	Dist(j-3,t+1)
bProb(j,k(t+1))	bProb(j-1,k(t+1))	bProb(j-2,k(t+1))	bProb(j-3,k(t+1))
+			
MIN $\left[\begin{matrix} \text{Dist}(j,t) + a\text{Prob}(j,j) \\ \text{Dist}(j-1,t) + a\text{Prob}(j,j-1) \\ \text{Dist}(j-2,t) + a\text{Prob}(j,j-2) \end{matrix} \right]$	MIN $\left[\begin{matrix} \text{Dist}(j-1,t) + a\text{Prob}(j-1,j-1) \\ \text{Dist}(j-2,t) + a\text{Prob}(j-1,j-2) \\ \text{Dist}(j-3,t) + a\text{Prob}(j-1,j-3) \end{matrix} \right]$	MIN $\left[\begin{matrix} \text{Dist}(j-2,t) + a\text{Prob}(j-2,j-2) \\ \text{Dist}(j-3,t) + a\text{Prob}(j-2,j-3) \\ \text{Dist}(j-4,t) + a\text{Prob}(j-2,j-4) \end{matrix} \right]$	MIN $\left[\begin{matrix} \text{Dist}(j-3,t) + a\text{Prob}(j-3,j-3) \\ \text{Dist}(j-4,t) + a\text{Prob}(j-3,j-4) \\ \text{Dist}(j-5,t) + a\text{Prob}(j-3,j-5) \end{matrix} \right]$

Figure 3. The recursion equation in SIMD form. Notice the left-most column is the scalar version of the equation.

The Internet Streaming SIMD Extensions can be used to process four distances at a time as seen in figure 2. There are two ways the Streaming SIMD Extensions can be used to implement the Viterbi decoding algorithm. The algorithm can be implemented using 16-bit integer SIMD or single-precision floating-point SIMD. The integer SIMD implementation provides the best performance and the single-precision floating-point SIMD can be used if a greater dynamic range (or precision) is needed. The following code demonstrates how to process four distances at a time with integer SIMD:

```
FourStateLoop:
movq mm0, [edi+eax] ; mm0 = [d3 d2 d1 d0 ]
movq mm1, [edi+eax+2] ; mm1 = [d4 d3 d2 d1 ]
paddsw mm0, [ebx+eax] ; mm0 = [d3+a3 d2+a2 d1+a1 d0+a0]
paddsw mm1, [ebx+eax+8] ; mm1 = [d4+a7 d3+a6 d2+a5 d1+a4]
movq mm2, [edi+eax+4] ; mm2 = [d5 d4 d3 d2 ]
paddsw mm2, [ebx+eax+16] ; mm2 = [d5+a11 d4+a10 d3+a9 d2+a8]
pminsw mm0, mm1
pminsw mm0, mm2 ; mm0 = min(Dist+aProb) = min(mm0,mm1,mm2)
paddsw mm0, [esi+eax] ; mm0 += bProb[obsVect[ecx]][i]
movq [edi+eax], mm0 ; store distance
add eax, 8
add ebx, 16
cmp eax, edx
jnz FourStateLoop

/* Process the remaining states if any (1, 2 or 3 remaining states) */
// ...
```

Figure 4. Processing four distances with integer SIMD.

Processing four data elements at a time and using the Internet Streaming SIMD Extensions like the `pminsw` instruction provide a significant performance gain in evaluating HMMs. For a detailed description on how to implement the Viterbi decoding algorithm in integer SIMD and single-precision floating-point SIMD, the reader is referred to the Intel Application Note AP-811 (see references below).

The Pentium III Internet Streaming SIMD Extensions can provide a significant performance gain in evaluating HMMs with the Viterbi algorithm. These new SIMD instructions can be used in many time-intensive applications if parallelism in the code exists. Speech recognition applications that use HMMs and the Viterbi decoding algorithm can utilize these instructions to provide faster recognition.

About the Author

Antonio C. Valles started with Intel's Software Performance Lab in 1997. Since coming to Intel, Antonio has been involved in performance tuning 2D/3D graphic engines, speech recognition, imaging, and MPEG Video applications. Antonio earned his Bachelor of Science Degree in Electrical Engineering from Arizona State University in 1997.

For More Information

- Speech Recognition: No Longer a Dream but Still a Challenge, Quinzel, Richard A., EDN, January 19, 1995, pgs 41–46.
- A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition, Rabiner, Lawrence R., Proceedings of the IEEE, February 1989, Vol. 77, No. 2.
- [Hidden Markov Model with Viterbi Decoding](#), Intel Application Note, AP-811, Copyright 1999.

Top Stories

Intel® 752 Graphics Accelerator: The Complete Multimedia Solution

By Bob Gregory
Director of Marketing
Graphics Components Division
Intel Corporation

Optimized for the Intel® Pentium® III processor, the Intel® 752 graphics accelerator delivers breakthrough visual quality, enhanced 3D rendering, advanced digital video performance, flexible display support and an upgrade path to AGP4X.

Compelling visual quality is important to consumer and business PC users alike, and that makes the quality of the graphics controller an important differentiation tool for PC OEMs. The Internet is creating new demands for high-performance video, 2D, and 3D visualization through entertainment and e-commerce applications. To date, the task of choosing the right graphics controller has been far from simple for developers. Technology advances quickly. Quality drivers and hardware can be hard to come by. In this fragmented environment, many graphics controller manufacturers resort to adapting last year's enthusiast-level controller to the mainstream volume PC segment.

It is no wonder that buyers in the mainstream PC segment often experience wide variations in graphics-related product quality. To help the industry address the issues of quality and compatibility, Intel is now complementing the processor-based PC segmentation model by introducing a graphics controller optimized for each PC market segment.

Intel is launching this effort by introducing its two next-generation graphics technology, built upon the Intel® Scalable Graphics Architecture (ISGA). This architecture provides a common graphics instruction set, implemented in both discrete and integrated products. ISGA will deliver software compatibility across a range of graphics products, providing driver and application stability. The first products in the new family are the Intel® 752 graphics accelerator, optimized for mainstream PCs, and the Intel® 810 chipset with integrated graphics, optimized for Value PCs.

Both the Intel 752 graphics accelerator and Intel 810 chipset build on the success of the award-winning Intel740™ graphics accelerator, which introduced a dramatic new level of graphics realism and quality in volume PCs.

It is important to note that while they are aimed at different user needs, Intel graphics controllers with ISGA will have compatible drivers and a scalable architecture both aimed at delivering a new level of quality, completeness, and stability for the industry. The ultimate goal is to reduce costs for OEMs and to provide consistent, high-quality choices for IT managers and end users.

Enhanced Performance, Plus Intel® Upgradeable AGP

Optimized for the Intel® Pentium® III processor, the Intel 752 graphics accelerator delivers visual quality for today's demanding consumer and business software applications. Backed by Intel's quality drivers and hardware, it is a feature-rich, high performance and versatile multimedia accelerator that delivers enhanced 3D rendering, advanced digital video performance, support for leading flat-panel displays, and an industry first: Intel® Upgradeable AGP.

Advanced technology, high-performance, and future upgradeability combine to make the Intel 752 graphics accelerator the complete multimedia accelerator for volume PC graphics. Here are the highlights:

- Breakthrough visual quality advances the state-of-the-art of 2D, 3D, and video performance. The pixel precise engine extends Intel's hyperpipelined architecture with new features, such as a 16 tap anisotropic filter, emboss bump mapping, texture compression, and texture compositing.
- Advanced digital video technology completes the multimedia experience and supports the convergence of PCs with digital video streams. Precision hardware motion compensation enables exceptional DVD playback. Flexible support for video capture, TV-in, and MPEG-2* video streams create a dynamic multimedia experience.
- Leading flat-panel solution that supports revision 1.0 of the Digital Visual Interface (DVI) specification, authored by Intel and the Digital Display Working Group (DDWG). This advanced digital connectivity specification protects investment in current display technology, while allowing OEMs to easily take advantage of new innovations that improve the visual computing experience.
- Intel Upgradeable AGP is an industry first, providing a simple and cost-effective migration path from AGP2X to AGP4X. This technology enables an Intel 752 graphics accelerator down on the motherboard with an AGP4X slot for future upgrades to the platform.

The Best Graphics Accelerator for the Internet

The Intel 752 graphics accelerator is optimized around the Pentium III processor streaming SIMD instructions, making it quite simply the best graphics controller for the Pentium III platform. It is validated with Intel® WebOutfittersm tools to ensure that leading-edge Internet, e-commerce, and entertainment applications are seen at their best. The Intel 752 graphics accelerator features consistent color rendering technology created through sub-color precision, and independent color and gamma controls to deliver accurate and high-impact visualization of Web content.

Simplifying Choices for OEMs

Quality Intel® drivers and the Software Developer's Kit (SDK) help OEMs simplify development with the Intel 752 graphics accelerator, while reducing development costs. Designed especially for the needs of today's mainstream performance PC users, the Intel 752 graphics accelerator provides the perfect balance of 2D, 3D, and video, and offers an easy migration path to the next key industry initiative: AGP4X. For enterprise business graphics applications, the accelerator supports Intel's Wired for Management technology to reduce total cost of ownership.

In the process of balancing performance, cost, and consumer appeal, PC OEMs face complex and difficult choices. Choosing the right graphics accelerator just became a lot simpler. Now is the time to start designing with this versatile solution.

For More Information

Visit the Intel Developer Web site for the [product overview and technical documentation](#). Read this month's focus story on [Intel Graphics Technology: Your Window to the Web](#).

About the Author

Bob Gregory has been director of marketing for the Intel Graphics Components Division since its inception in 1998. He previously served as director of strategic planning for the Intel Platform Components Division, where he was responsible for new product planning and participating in industry initiatives. Prior to joining Intel, Bob held management positions in product marketing and product planning for desktop products at Compaq Computer Corporation. Before joining Compaq, he worked at Panasonic Company as a product manager for PC products. He has a master's degree in business administration from Fairleigh Dickinson University and a bachelor's degree in electrical engineering from Villanova University.

Network Processors Help Enable the New Internet Economy

By Doug Carrigan
Product Marketing Engineer,
Networks Communications Group
Intel Corporation

Network processors are software-programmable silicon engines that enable network systems vendors to quickly incorporate new technologies and deliver enhanced services, while saving time-to-market and development costs.

The new Internet economy will be built on a foundation of secure transactions which cross all the domains of corporate and personal computing. This will require the integration of business critical data services with support for voice and video transmission all in a single, integrated network. These integrated networks will be dependent upon the rapid deployment of new technologies to support both Local Area Network (LAN) and Wide Area Network (WAN) traffic while simultaneously providing advances in security, Quality of Service (QoS), and manageability. These networks can only be built on equipment which supports rapidly evolving features, many of which are based on standards efforts that takes years to come to fruition.

The convergence of voice and data networks has led to an escalation of the competitive pressures facing equipment vendors on both sides of the voice and data networking collision. Two competitive markets have been thrown together practically overnight and billions of dollars of equipment and services revenue is up for grabs. The competition in this segment of the market has helped drive all of the major players to increase the rate at which they deploy new products.

The drive to integrate multiple disparate networks has created an environment where network and communication system vendors have begun to re-evaluate the way they design products. New design paradigms are being explored in an attempt to provide more flexibility to adapt to changing market requirements and evolving standards and to extend the useful life of deployed products. It is this environment that has brought about the emergence of a new class of silicon product—the Network Processor.

The Network Processor is designed to provide the data movement performance that has historically been provided by dedicated Application Specific Integrated Circuit (ASIC) solutions. Its key benefit, however, will be its flexibility to support the sophisticated packet filtering and forwarding capabilities demanded by integrated voice/data networks, particularly at the LAN/WAN edge. This flexibility helps the systems vendor to develop and deliver differentiated products. It is the combination of high performance, flexibility, and the ability to support OEM differentiation that will set the Network Processor apart from other solutions.

Network Processors: Defined

Network Processors are software-programmable, silicon engines designed to provide a wide-range of enhanced features and capabilities across diverse communications platforms. Network Processors are designed to offer, in a programmable engine, the advanced switching and routing capabilities normally designed in complex ASICs. Network Processors are also designed to combine the real-time control and management capabilities of an embedded processor to support scalable system designs.

The advanced capabilities offered by Network Processors include:

- Multi-protocol switching to enable integrated LAN/WAN switching systems and support legacy protocols
- Multi-layer programmable switching and routing capabilities to support sophisticated, policy-based traffic management
- Programmable Quality of Service support to enable traffic prioritization across disparate networks

In addition to programmability, Network Processors must also deliver the “wire speed” performance expected of enterprise and carrier-class networking systems and offer continuous, predictable performance gains over time.

Network Processors Help Speed Time-to-Market and Reduce Risk and Development Cost

The capabilities provided by the Network Processor help network systems vendors to more quickly incorporate new technologies and deliver enhanced services. Network Processors provide a platform upon which advanced features may be incrementally built. By reducing the need to “reinvent the wheel” through repeat development and verification of existing features, Network Processors provide reductions in time-to-market and development cost. The flexibility programmable Network Processors offer supports system vendors by letting them shift development resources to areas that add value and increase product differentiation.

Developers are frequently required to begin projects by anticipating market requirements and making feature set decisions early in the process. The risk is that the feature sets they choose may not be correct. In these cases, the project must be either restarted or delivered without vital features.

The Network Processor alleviates this risk by providing a standard platform that lets feature set definition continue throughout the development process. Removing dependence on hardware availability also allows concurrent hardware and software design, so Network Processor-based systems can be deployed while additional features are still in development. In systems that have already been deployed, Network Processors provide the ability to implement field upgrades. Thanks to this software and silicon solution, hardware obsolescence is significantly reduced.

A Not-So-Hypothetical Example

To illustrate how easily Network Processor benefits can be realized, let's examine the case of a communication system vendor designing equipment to support integrated voice and data communications. The switching system must be capable of supporting the Quality of Service required to deliver high-quality voice transmission across both LAN and WAN circuits.

Lack of a standard mechanism for communicating the required priority from one network type to another means that the system vendor must implement a proprietary scheme and risk being incompatible with an eventual standard. The alternative is to take the risk of developing a highly complex ASIC which offers sufficient flexibility to support all possible implementations of the standard.

The Network Processor offers a more attractive alternative: implement the traffic prioritization mechanism in the Network Processor. The system vendor is then able to deploy the product in advance of the final standard yet feel confident that the product can be reprogrammed in the field to be standard-compliant. The same technique can be used to deploy proprietary features which can be “turned on” when connected to another of the same vendor's products and default to a standard-compliant mode to guarantee interoperability with another vendors product.

Network Processors: A Key Component of Intel's Internet Silicon Building Block Strategy

Network Processor technology represents a key element in Intel's strategy to be a premier building block supplier to the new Internet economy. It is a key enabler to accelerate the deployment of manageable, scalable, intelligent networks which are central to this new economy. Further, the Network Processor is a logical evolution of Intel's network and communications silicon business—applying Intel's expertise in developing and delivering processor and networking silicon to network and communications system vendors.

Network Processors will be the centerpiece of a large and growing system silicon portfolio, joining Intel's traditional Ethernet controller, physical layer, and system silicon products. Intel has demonstrated its commitment to being a leader in this space with its announcement to acquire Level One Communications. The Level One acquisition is expected to add breadth and depth to Intel's network and communications silicon development and support capabilities.

Intel has committed resources to help deliver and support Network Processors as standard products to the industry at large over the long term. This will provide the system vendor the security necessary to adopt the Network Processor as a critical element of their design methodology. Intel will continue to build upon its strength in the network and communications segments to help in the deployment of the Network Processor.

About the Author

Doug Carrigan is a product marketing engineer in Intel's Network and Communications Group working on Network Processors. Doug has over 15 years of experience in the development of embedded microprocessors and networking silicon products. Doug joined Intel last year from Digital Equipment Corporation.

For More Information

Please visit us regularly for updates on our [Network Processor efforts](#). More information on Intel's Network Processor plans can be found at the [Intel pressroom](#).

Self-Test Tool Accelerates WfM 2.0 Implementation

By José Ramirez
Senior Software Engineer
Intel Architecture Labs
Intel Corporation

The new Wired for Management 2.0 Self-Test Toolkit from Intel provides a straight-forward testing mechanism to ensure that WfM 2.0 specification technologies are functioning correctly.

There's no question that the WfM 2.0 specification will have a significant influence on improving the manageability of connected PCs across the enterprise. By adding advanced manageability capabilities for desktop, mobile, and server systems to the existing specification, WfM 2.0 is poised to help companies gain control of their computing environments, increase their business flexibility, and reduce total cost of ownership (TCO) of their computing resources. WfM v2.0 enhances basic management capabilities such as system component instrumentation, universal network boot, power management, and remote wake-up while adding streamlined help desk support capabilities.

One of the keys that will help unlock the door to the advantages that WfM 2.0 provides for platform developers as they migrate to the new specification is a robust self-test toolkit. But while there has been a limited distribution of beta versions of the self-test tools, nothing has been publicly available across the industry to help developers ensure that their initial implementations of WfM 2.0 technologies are functionally correct. All that's about to change. On April 30, just in time to prepare for the [WfM 2.0 plugfest](#), Intel will make the WfM 2.0 Self-Test Toolkit available for downloading—free of charge—from the Intel Developers' Web site.

A Suite of Seven

The new self-test toolkit consists of tool documentation and seven different tools designed to help accelerate the implementation of WfM 2.0 technologies:

- **DMI Compatibility and Instrumentation Checker** tests compliance with the Desktop Management Interface V2.0 standard.
- **Remote Wake-up Utility** tests support for the five sleep states (S1-S5) of the ACPI standard (support for at least one of the sleep states is required in the WfM 2.0 baseline specification). The specification extends the benefits of the Instantly Available PC (IAPC) to include managed mobile computing platforms.
- **PXE Server/Client Utilities** provide self-test capabilities related to WfM 2.0's improved architecture for the Pre-boot eXecution Environment.
- **SM BIOS Test Tool** provided by Dell Computer helps developers verify proper operation of the latest version of the SM BIOS and WfM 2.0 specifications, which jointly provide a standard method for accessing a system's platform data.
- **SNMP Compatibility Tool**, similar to the DMI checker, checks SNMP traps and DMI to SNMP mappings.
- **Boot Integrity Services (BIS) Test Tool** tests for the proper downloading of boot images to be authenticated using digital signature technology for BIS-enabled systems.
- **Problem Resolution Software (PRS) Test Tool** is available for testing PRS, which implements an agent used by WfM-enabled systems to share support knowledge and exchange trouble tickets.

For a desktop system to be identified as WfM 2.0 enabled, it is required to pass the tests of these four tools: the DMI checker, Remote Wake-up, PXE, and SM BIOS (these tools test the required features of the WfM 2.0 specification). The remaining three tools, SNMP, BIS, and PRS, are for testing features that are currently recommended—but not required—by the WfM 2.0 specification. A user-friendly installation program enables developers to easily install only the tools they need to test their specific systems.

Easy Self-Test

To facilitate testing, the following is offered as a recommended test environment:

- A Windows NT* 4.0 server that includes Service Pack 3 with SNMP and DHCP services installed and configured correctly. The server should also have an installed DMI 2.0 service provider and a DMI browser such as DMI Explorer included with LANDesk® Client Manager or DCTS—both available from Intel.
- A second server, which only requires that the PXE PDK be installed and configured, to test server arbitration.
- The system under test should have an operating system that supports plug and play for testing USB and other hot-swappable devices; for example, Windows* 98 supports USB, Windows NT does not.

Wired for Management has proven to be one of the most influential and widely adopted initiatives in the PC industry over the last few years, and WfM 2.0 promises to extend its reach even further. With the WfM 2.0 Plugfest on the near horizon, developers would be well-advised to download and use the new WfM 2.0 Self-Test Toolkit starting on April 30, in order to best prepare for the Plugfest and accelerate the implementation of the latest WfM capabilities in their next-generation products.

About the Author

José Ramirez is a senior software engineer for the Intel Architecture Labs (IAL), where his responsibilities include developing WfM tools and supporting the WfM V2.0 Self-Test Toolkit.

For More Information

Visit these sites for more information:

- [Wired for Management](#)
- [WfM 2.0 Plugfest](#)

The Benefits of Integrated RAID

By Christopher C. Croteau
Marketing Programs Manager
I/O Products Division
Intel Corporation

The Internet, electronic commerce, and other business-critical applications are driving a rapid growth in storage technology. The use of Intel® i960® Intelligent I/O processors in RAID controllers builds a strong foundation for data protection and server performance.

The growth of the Internet, electronic commerce, and other business-critical applications are spurring a rapid growth in storage. Businesses of all sizes face the challenge of how to protect data that can literally represent the future of the enterprise. When a business relies on servers, even minor unplanned disruptions to data access can have immediate consequences for the company's productivity and revenue stream. Protecting disk storage with RAID (redundant array of independent drives) technology is therefore a vital issue for any server-based business. The use of Intel® i960® Intelligent I/O processors in RAID controllers builds a strong foundation for data protection as well as optimal server performance.

RAID will increasingly be available as an integrated feature of entry-level and midrange Intel® Architecture (IA) servers. In addition, Integrated RAID on the motherboard is part of a widespread effort to boost server performance and set the stage for the adoption of new server I/O technology initiatives.

Opportunity for Developers

Aberdeen Group research indicates that fewer than half the servers in the sub-\$8,000 price band currently feature RAID subsystems. This is at a time when business users are increasingly focusing their attention on server reliability and minimizing system downtime even as they continue to deploy IA servers for business-critical applications. Developers have a great opportunity to meet this emerging need by adding Integrated RAID to their systems.

Why Hardware-Based RAID is the Preferred Solution

While RAID functions can be handled by the server's CPU (host-based RAID), this approach can significantly degrade overall server performance. This is due to the fact that implementing RAID functions in main memory instead of on an Intelligent I/O processor-based controller can consume valuable CPU cycles in addition to taking up space in the server's main memory.

The CPU burden can be significant. "Host-based RAID Level 0 will use about 20 percent of the available CPU cycles," says Mylex Corporation's Director of Worldwide Channel Marketing Mike Joyce. The solution to this problem is the integration of RAID functions with Intelligent I/O processors between the server's CPU and the disks themselves. The dedicated I/O processor can handle many I/O tasks, such as interrupt processing and RAID parity generation, freeing host CPU cycles for application processing.

The Benefits of Integrated RAID

Integrated RAID on the motherboard enables the introduction of RAID technologies to entry-level servers. It can alleviate the cost and deployment concerns that may have delayed adoption of the technology by many business users. In addition, adding RAID as part of a motherboard architecture offers developers a more cost-effective and easy-to-deploy RAID solution, especially for entry-level IA servers.

In its white paper entitled *Integrated RAID: Building More Storage Functionality into Servers*, the Aberdeen Group identifies seven key benefits of integrated Intelligent RAID:

- Boosting availability.
- Easier deployment of storage clusters.
- Better use of server resources to handle multiple applications.
- Reducing total system costs.
- Improved manageability.
- Building a foundation for future technologies that can build on the system's RAID environment.
- Offering users at all levels a way to increase server availability and reliability.

The final point has important significance for server OEMs and developers. As reliability continues to become a focal point for MIS in businesses of all sizes, integrated RAID will become a key differentiator.

Building a Foundation for Next-Generation Technologies

The use of dedicated Intelligent I/O processors will eventually enable devices attached to a server to communicate with each other by means of peer-to-peer communications without involving the server's main CPU. Intelligent I/O processors will also assume an expanded role in overall I/O tasks, giving the server I/O system the role of a separate intelligent device from the host processor.

Server OEMs and developers should take advantage of the opportunity to deliver Integrated RAID in their entry-level and midrange systems. The integration of RAID inside IA servers is playing a growing role in boosting server I/O system performance. RAID controllers based on Intelligent I/O processors provide the foundation for future applications that improve data access, bus bandwidth, and enterprise network performance.

For More Information

Read the [Aberdeen Group white paper](#) *Integrated RAID: Building More Storage Functionality into Servers*.

See how Mylex Corporation demonstrated the effect of [Intelligent RAID on workstation performance](#) at the Intel Developer Forum.

About the Author

Christopher Croteau is marketing programs manager for Intel's I/O Products Division, where he is responsible managing strategic and business programs. He has developed programs to promote the adoption of new technologies, including I2O* software architecture and USB.

Power Budgeting Software Protects the Instantly Available PC

By Patrick Bohart
IAPC Technical Manager
Intel Corporation

By Ram Chary
Staff Software Engineer
Intel Corporation

Intel's new Power Budgeting software application ensures that standby current capacities won't be exceeded on the motherboard when Instantly Available PCs enter the ACPI S3 sleep state.

Intel's efforts to advance the frontiers of Instantly Available PC (IAPC) technology in recent years are now being realized in the development of next-generation systems slated for delivery to market in the second half of 1999. Instantly Available PC technology allows OEMs to design full-featured, high-performance systems that meet stringent sleep-state targets—while delivering new, useful, and exciting platform capabilities.

IAPC employs a configuration-independent mechanism that removes power from all major subsystems—such as graphics, audio, the processor, and the hard drive—when asleep. This enables OEMs to continue to innovate in ways that deliver higher performance and new features without having to worry about regulatory impacts. The IAPC architecture also allows certain parts of the system, such as modems and network interface cards, to draw a small trickle current that enables them to wake the machine when an external event occurs.

The Current Situation

While IAPC technology promises to have a significant effect on the next generation of PC platform systems now in development, one potential problem looming on the horizon concerns the standby current capacities of today's power supplies. That's because the PC industry is now using a standard ATX power supply that typically only supports a maximum of 720 mA of standby current. This current is used to power-up a system's ACPI controller, SDRAM and the 3.3 Vaux and PME# pins on the motherboard when the system enters the S3 sleep state, as specified by IAPC technology in the [Instantly Available Power Managed Desktop PC Design Guide](#).

Given the standard configuration of most motherboards today, 720 mA may not supply enough current to handle heavy loading situations—as in the case of systems that employ a multitude of wake-enabled add-in cards attached to the PCI bus (each potentially capable of consuming up to 375 mA). Exceeding this current capacity could cause the system to fail on resume. And, that's where Intel's new Power Budgeting software can help.

On the Lookout

Developed by the Intel Architecture Labs (IAL), Power Budgeting is an application that scans PCI bus configuration space to determine the amount of standby current that will be required by each add-in card while the system is in suspend-to-RAM (S3). Functioning in much the same way as an anti-virus program, the application loads up in the system launch and references the combined current of all power-managed devices on the PCI bus against the known 5V standby current capacity of the power supply. If there's a discrepancy, the application automatically and transparently transitions the maximum sleep state—from S3 to S1—a level that the power supply can handle.

In essence, Power Budgeting acts as an effective insurance policy for platform developers, enabling them to implement IAPC technology in their systems without fear of potential problems associated with standby current overload on the motherboard. It's clear that the industry will eventually move to more robust power supplies with higher standby current capacities, particularly as USB wake-up continues to emerge. But until that happens, Power Budgeting should prove to be an ideal software for eliminating what otherwise could have become a thorny issue for platform developers.

An Eye Toward the Future

Power Budgeting is available today for OEMs implementing IAPC technology. As such, it represents part of Intel's ongoing efforts to support IAPC, an important component of the Easy PC Initiative for promoting ease-of-use on the PC platform moving forward. Intel is actively looking for beta testers for the Power Budgeting software, and encourages developers to contact the company for more information.

Where do we go from here? More robust power supplies will be in order, as will integration of IAPC and associated technologies into Windows* 2000. In addition, independent hardware vendors will also be charged with the task of trimming standby current consumption in their succeeding generations of add-in card products. And a steady migration toward USB is inevitable. For that reason, Intel is developing a version of Power Budgeting that will support USB, to complement the current PCI bus support provided by the software.

About the Authors

Patrick Bohart is the initiative technical manager for Intel's Instantly Available PC technology. His responsibilities include providing technical direction, driving compatibility standards, and overseeing communications activities related to Intel's IAPC enabling efforts.

Ram Chary is a staff software engineer in the Platform Architecture Lab. His focus is on enabling the deployment of Instantly Available PCs. During his 14 years at Intel, Ram has developed processor microcode, real-time kernels and networking products as well as managed the engineering teams for multiple ProShare product releases.

For More Information

Check out the [Instantly Available Technology Web site](#) for more information on Power Budgeting and IAPC Technology.

Technology News Bytes

April 20

Discount on Test Drive for Developers

Intel and RealNetworks are offering developers a discounted test drive of RealSystem* G2, with Intel® Streaming Web Video technology. Features of this first-time offer‡ include:

- RealProducer* Plus*, a top content creation tool.
- RealDeveloper* Program membership.
- Three months of easy content hosting with RealBroadcast Network.*

This offer is good until June 30, 1999. For more details on RealSystem G2, check out the [Real](#) Web site.

‡This offer is subject to conditions and limitations set by RealNetworks. See www.real.com for more information.

April 19

Intel and NBC Work to Speed the Development of Enhanced DTV

Intel and NBC today announced at the National Association of Broadcasters (NAB) convention that the companies would collaborate on a multi-year project to create and provide enhanced digital television programming. NBC's enhanced DTV programming is scheduled to be available to viewers in the fall of 1999. NBC's enhanced DTV programming will provide a variety of new experiences for viewers on multiple digital platforms.

April 19

Intel to Enter Network Processor Market Segment

Intel announced that it will offer network processor components designed to enable new features and services to be added quickly and economically to multi-protocol enterprise- and service provider-class switches, routers and access concentrators. These software-programmable forwarding and control engines will be designed to help data and telecommunications equipment vendors to more quickly incorporate rapidly evolving technologies and reduce development costs, risk and time to market.

April 15

Intel and Metacreations Announce Availability of Specification for Scaling and Streaming 3-D Content

Intel MetaCreations Corporation announced at Spring Internet World the immediate availability of a specification for scaling and streaming 3-D content on the Web, the MetaStream* 3-D file format designed for the Intel® Architecture. The open specification, developed by Intel Architecture Labs and MetaCreations, will enable Web application developers to obtain the benefits of the MetaStream format. The companies intend to jointly continue to develop the capabilities of the MetaStream 3-D file format and companion tools, such as the MetaFlash* 3-D camera being brought to market later this year by Kodak and Minolta and MetaCreations' newly announced Canoma* software.

April 14

Intel Showcases New Technology at Spring Internet World '99

Spring Internet World '99, promoted as the world's largest e-business and Internet technology event, took place April 14-16 at the Los Angeles Convention Center. Intel showcased an array of powerful technologies and business solutions, including solutions based on the new Pentium® III and Pentium® III Xeon™ processors.

April 14

Intel Senior Vice President Sean Maloney Kicks off Spring Internet World

Intel is enabling further innovation in the Internet economy. During his keynote speech, Intel Senior Vice President Sean Maloney provided an overview of Intel's role in this market segment. The role Intel's chips play in the PCs, servers and networking equipment that make up the Internet, and the equity investments the company makes to fuel the growth of E-commerce technologies, content and services were some of his topics.

April 13

[Intel First Quarter Revenue \\$7.1 Billion](#)

Intel announced first quarter revenue of \$7.1 billion and earnings of \$2.0 billion or \$0.57 per share, adjusted for the 2-for-1 stock split paid on April 11, 1999. Early in the second quarter, Intel shipped its 100 millionth microprocessor based on the P6 microarchitecture.

April 12

[Intel Provides Small Businesses with Faster Internet Access](#)

Intel announced the addition of analog modem bonding to the Intel® InBusiness™ Internet Station helping to significantly increase availability and speed for small business customers connecting to the Internet. Whereas an average analog Internet connection transfers information, faxes, or email at up to 56 Kbps, analog modem bonding combines the capability of up to three modems for approximately 168 Kbps, a speed faster than that of ISDN.

April 12

[Intel and Excite to Develop New Personalized E-Commerce Service](#)

Intel and Excite, Inc. announced a multi-year agreement to develop a new e-commerce service designed to simplify buying and selling on the Internet. The service will combine the strengths of Intel's commerce technologies with Excite's personalization, direct marketing expertise and reach. Together the two companies will build a service that provides an easy way for consumers to locate and purchase a very broad array of products and services tailored to their shopping interests.

April 8

[ATI Announces Support for Intel's Digital Visual Interface](#)

A key technology enabler for the Digital Display Workings Group's new Digital Visual Interface (DVI), ATI Technologies Inc. announced that it will develop graphics products for the digital flat panel (DFP) market based DVI standard developed by the Intel Corporation-led Digital Display Working Group (DDWG). DDWG today announced the release of Revision 1.0 of the DVI specification. The Digital Visual Interface was introduced by Intel and the DDWG in August 1998, and the first specification of the new interface was presented in February 1999.

April 7

[Industry Support Grows For Intel and Microsoft Easy PC Initiative](#)

At the Windows® Hardware Engineering Conference and Exhibition (WinHEC) 99, Microsoft Corp. President Steve Ballmer outlined the Easy PC Initiative, a multiyear vision co-developed with Intel Corp. in continuation of work on PC ease of use and targeted at improving the overall experience for PC users. Leading PC manufacturers, including Compaq Computer Corp., Dell Computer Corp., Fujitsu Microelectronics Inc., Gateway 2000 Inc., Hewlett-Packard Co., IBM Corp., Micron Electronics Inc., NEC USA Inc. and Toshiba Corp., announced their support for the initiative, and will work to develop new PC prototypes by the end of this year.

April 6

[Intel Brings Powerful, Simple PC Networking to the Home](#)

Intel introduced the AnyPoint™ Home Network product line, which helps make it easy for families with more than one PC to share Internet access, printers, files and games. The products will be available at intel.com, CompUSA and Gateway. The AnyPoint Home Network uses existing phone lines to connect home PCs without the need for additional wiring. The AnyPoint Home Network is the first phoneline solution that can be installed without opening the PC chassis, because it attaches to the parallel port of home PCs.

April 6

[Intel Introduces New Stackable Hubs and Standalone Switches](#)

Intel announced the Intel® Express 330T Stackable Hubs and the Intel® Express 410T Standalone Switches, furthering its commitment to deliver feature-rich products for workgroup and desktop connectivity. The Intel Express products offer increased flexibility for evolving networks with intuitive network control and configuration at a compelling price.

April 5

Intel Introduces New Mobile Celeron™ Processor 333MHz for Low-Cost Mobile PCs

Further enhancing its role as the leading supplier of processors for the value PC market segment, Intel introduced the mobile Intel® Celeron™ processor at 333MHz. The new mobile Intel Celeron processor at 333MHz offers consumers and businesses the benefits of mobility along with great performance at an exceptional price. Intel's mobile Celeron processors offer the benefits of low power for long battery life, unique packaging for smaller and lighter systems, and a mobile-friendly thermal design specification.

March 31

Intel® StrongARM®* Processor, Companion Chip Optimized For Handheld Computing Devices

Intel announced the addition of the Intel® StrongARM®* SA-1110 processor and the SA-1111 companion chip to the Intel StrongARM product portfolio. The two provide increased memory and I/O design flexibility for handheld product manufacturers, and complement the performance, power, and cost benefits of the current SA-1100 processor and SA-1101 companion chip. The building blocks for a range of handheld applications, the SA-1110 processor enables new capabilities in handheld computing devices such as software-based communications, improved speech recognition, and better handwriting recognition.

March 30

Remedy, Intel and Symantec Integrate Industry-Leading Solutions to Create IT Service Management Platform

Remedy® Corporation, Intel, and Symantec Corporation announced that the three companies have joined forces to provide IT organizations a standards-based, integrated platform for IT Service Management. The relationship builds on existing agreements between the companies, which have been working together for more than a year to deliver increased integration to their customers.

March 29

Micron Electronics Selects Intel® AnswerExpressSM Support Suite as Value-Added Option on Consumer and Small Business Desktop PCs

Micron Electronics, Inc., a leading direct vendor of personal computer systems, announced that it has selected Intel® AnswerExpressSM Support Suite as a value-added option for Micron's small business and consumer customers. Intel AnswerExpress Support Suite is an online support service that provides consumers and small businesses with live PC support, virus protection, online backup and an answer library.

March 29

Intel Corporation Announces Low Voltage, 1.8 Volt Advanced+ Boot Block Flash Memory

Intel announced the 1.8 Volt Intel® Advanced+ Boot Block, providing low voltage and fraud protection capabilities. These features are key to battery-powered handheld devices such as cellular phones, pagers, PC companions and Global Positioning Satellite (GPS) systems. The new Intel 1.8 volt Advanced+ Boot Block flash offers 90 ns read access and uses up to 25 percent less energy than similar 3 volt products.

March 26

The Intel® WebOutfitterSM Service Goes Live

Intel announced that the Intel® WebOutfitterSM service (intelweboutfitter.com), a new Internet service for Intel® Pentium® III processor owners, is now live on the web. The Intel WebOutfitter service delivers cutting-edge content, tools and tutorials designed to enhance the Internet experience for consumers of Pentium III processor-based computers. Intel plans to evolve this new service to include next-generation content, technologies and services.

Industry Events

May 1999

E3 (Electronic Entertainment Expo)

Conferences and Workshops 12–14

Exposition 5/13–5/15

Los Angeles Convention Center

Los Angeles, CA

Electronic Entertainment Expo (E3), billed as the world's largest annual tradeshow for the nation's fastest growing entertainment industry, brings to Los Angeles the latest and most advanced PC and video game hardware and software, and tens of thousands of the luminaries who make the decisions shaping the industry's destiny. E3 attendees from around the world will come to buy, sell, and invest in the interactive entertainment industry's headlining releases and groundbreaking technologies. E3 is owned by the Interactive Digital Software Association (IDSA), the U.S. association dedicated to serving the business and public affairs interests of companies that publish video and computer games for video consoles, personal computers, and the Internet.

Intel executive vice president and general manager, Intel Architecture Business Group, Paul Otellini, will be presenting a keynote on May 14.

IDUG 11th Annual North American Conference

May 16–20, Orlando, Florida

IDUG, the International DB2 Users Group, is an independent, not-for-profit, user-run organization whose mission is to support and strengthen the information systems community by providing the highest quality education and services designed to promote the effective utilization of the IBM DB2* Universal Database family of products. This conference in Orlando is the 11th annual IDUG North American Conference. It features over 120 technical sessions, along with SIGS, panels, and networking opportunities. The technical session agenda adds several sessions focusing on DB2 in the growing Enterprise Resource Planning arena. Other featured sessions for this year's conference include research and development updates in DB2 for OS/390*, DB2 for UNIX, Windows*, and OS/2*; Business Intelligence; and e-business.

Corporate Computing Expo 99

May 18–20

Moscone Convention Center

San Francisco, CA

CCX, Corporate Computing Expo, is an industry gathering for small- to mid-sized businesses. According to show literature, technology has forever altered the way you do business—making the size of your company an issue of the past. At CCX, you can put together the right technology solutions for your business model to stay on top of your market, increase your profits, and maintain your competitiveness. CCX has brought more than 100 of the country's leading software, hardware, and technology service providers together, each featuring state-of-the-art, practical small- and mid-sized business solutions. Intel will be demonstrating client-server solutions.

PC Developers' Conference and Expo

May 24–27

Santa Clara Convention Center

Santa Clara, CA

A must-attend event for design professionals and managers involved in designing PC-compatible embedded systems and applications, this show is co-sponsored by Intel Embedded Microcomputer Group. Intel's Jim Pappas will be presenting the conference keynote at 1:45 on Tuesday, May 25.

June 1999

Computex 99

June 1–4, Taipei, Taiwan

At this show for resellers and system integrators, look for Intel to demonstrate graphics and chipsets, including the new Intel® 810 chipset

SIA (Securities Industry Association)

June 15–17

New York Hilton Hotel

New York, NY

This is the annual management education program designed expressly for individuals within securities firms who have management responsibility for technology-based support activities, including telecommunications, market data, trading systems, processing and information systems. The three general sessions, 20 workshops, luncheon meeting and Exhibit are all aimed at helping the target audience become more effective at their firms; a very challenging goal. The Conference program is developed under the guidance of the sponsoring Technology Management Committee.

DAC (Design Automation Conference)

June 21–25

New Orleans, LA

With the theme of Designing Systems for the New Millennium, the Design Automation Conference (DAC) bills itself as the premier Electronic Design Automation (EDA) and Silicon solution event. The 36th Design Automation Conference is sponsored by the ACM/SIGDA (Association for Computing Machinery/Special Interest Group on Design Automation), IEEE/CAS (Institute of Electrical and Electronics Engineers/Circuits and Systems Society), and EDAC (Electronic Design Automation Consortium).

DAC features an outstanding technical conference with over 50 sessions led by leading system designers and researchers presenting the latest in design methodologies and EDA tool developments along with industry trends and information. The DAC Exhibition and Demo Suite area features over 230 of the leading EDA, silicon, and IP providers.

PCEXPO 99

June 22–24

Javits Convention Center

New York, NY

With the theme Technology for Business, PCEXPO claims it's not just an event—it's the event for business technology, the most influential corporate computing event in the country. From Windows NT* to NetWare*, client/server to Java*, PCEXPO showcases the latest technology for business, not just on the desktop but also throughout the enterprise. Intel will have a large presence at PCEXPO. Demonstrations will include the 810 chipset, mobile processors, and wireless technology.

System Test Implementers Forum Planning Session

June 24–25, 1999

San Diego, CA

All members are invited to help define System Test Interface Forum assertion coverage beyond PC 99.

—End of Platform Solutions Issue 19—